

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
REQUEST FOR FILING APPLICATION UNDER 37 CFR 1.53(b)
WITHOUT FILING FEE OR EXECUTED INVENTOR'S DECLARATION

Assistant Commissioner for Patents
Washington, D.C. 20530 U.S. PTO

Atty. Dkt. 2380-155

Date: October 16, 2000

Sir:

10/16/00

This is a request for filing a new PATENT APPLICATION under Rule 53(b) entitled:
FAULT DETECTION IN MULTI-PLANE SWITCH

without a filing fee and/or without an executed inventor's oath/declaration.

This application is made by the below identified inventor(s). Attached hereto are the following papers:

- ☒ An abstract together with
39 pages of specification and claims including
20 numbered claims and also attached is/are
13 sheets of accompanying drawings.

☐ This application is based on the following prior foreign application(s):

Application No.

Country

Filing Date

respectively, the entire content of which is hereby incorporated by reference in this application, and priority is hereby claimed therefrom.

☐ This application is based on the following prior provisional application(s):

Application No.

Filing Date

respectively, the entire content of which is hereby incorporated by reference in this application, and priority is hereby claimed therefrom.

☐ Certified copy/ies of foreign applications attached.

☐ This application is a ☐ continuation/☐ division/☐ continuation-in-part of application Serial No. , filed
☐ Please amend the specification by inserting before the first line: --This application is a ☐ continuation/☐ division/
☐ continuation-in-part of application Serial No. , filed , the entire content of which is hereby incorporated by
reference in this application.--

☐ Please amend the specification by inserting before the first line: --This is a continuation of PCT application No.
 , filed , the entire content of which is hereby incorporated by reference in this application.--

☐ Please amend the specification by inserting before the first line: --This application claims the benefit of U.S.
Provisional Application No. , filed , the entire content of which is hereby incorporated by reference in this
application.--

☐ Preliminary amendment to claims (attached hereto), to be entered before calculation of the fee.
☒ Also attached. Information Disclosure Statement w/ PTO-1449

1. Inventor: Mats MI LINDBLOM Swedish
(first) (last) (citizenship)
Residence: (city) Johanneshov (state/country) Sweden
Post Office Address: Ahusvagen 6, Johanneshov, Sweden
(Zip Code) SE-121 51

2. Inventor: Hans MI BERGHAGER Swedish
(first) (last) (citizenship)
Residence: (city) Alvsjo (state/country) Sweden
Post Office Address: Slanbarstigen 23, Alvsjo, Sweden
(Zip Code) SE-125 56

NOTE: FOR ADDITIONAL INVENTORS, check box ☒ and attach sheet with same information.

Address all future communications to NIXON & VANDERHYE P.C., 1100 North Glebe Road, 8th Floor, Arlington,
Virginia 22201.

1100 North Glebe Road, 8th Floor
Arlington, Virginia 22201-4714

Telephone: (703) 816-4000

Facsimile: (703) 816-4100

HWB:lsb

NIXON & VANDERHYE P.C.

By Atty: H. Warren Burnam, Jr., Reg. No. 29,366

Signature: *H. Warren Burnam, Jr.*

jc326 U.S. PTO
09/688165
10/16/00

3. Inventor: Per SYDHOFF Swedish
(first) MI (last) (citizenship)
Residence: (city) Stockholm (state/country) Sweden
Post Office Address: Backvindeln 122, Stockholm, Sweden
(Zip Code) SE-129 42

Our Ref.: 2380-155
P12041US1

U.S. PATENT APPLICATION

Inventor(s): Mats LINDBLOM
Hans BERGHAGER
Per SYDHOFF

Invention: FAULT DETECTION IN MULTI-PLANE SWITCH

***NIXON & VANDERHYE P.C.
ATTORNEYS AT LAW
1100 NORTH GLEBE ROAD
8TH FLOOR
ARLINGTON, VIRGINIA 22201-4714
(703) 816-4000
Facsimile (703) 816-4100***

SPECIFICATION

FAULT DETECTION IN MULTI-PLANE SWITCH

BACKGROUND

This application is related to simultaneously filed United States Patent Application Serial Number 09/____,____ (attorney docket 2380-154), entitled
5 “SYNCRHONOUS CHANGE OF SWITCHPLANE”, which is incorporated herein by reference in its entirety.

1. FIELD OF THE INVENTION

The present invention pertains to cell switches, and particularly to cell switches which operate with plural switch planes.

2. RELATED ART AND OTHER CONSIDERATIONS

10 Cell switches, such as ATM cell switches or Internet Protocol (IP) routers, are extensively employed in communication systems, such as mobile telecommunication systems, for example. Typically a communication node with a cell switch has a switch core which routes an incoming cell from an ingress port of the core to one or more
15 egress ports of the core. The core ports are connected to respective switch port interface units (SPICs). In some systems, the switch port interface units reside on a board, with the board also having a board processor (BP) and (usually) other devices situated thereon. Some of the boards can be, for example, exchange terminals (ETs) which connect to links leading to other nodes of the network (e.g., other cell switches).
20 One of the boards may have a main processor (MP) which governs the operation of the cell switch node.

Usually, when a traffic cell is received at cell switching node it is received at an extension board. The header of the incoming traffic cell is analyzed and manipulated at the extension board in order to route at least the payload of the cell through the
25 switch. Various activities are performed at the switch port interface unit of the board,

including queuing of the traffic cell in an appropriate buffer of the switch port interface unit in accordance with its priority and destination switch port interface unit of the node. When the traffic cell is selected for dequeuing from its switch port interface unit, it is transmitted over a link to the corresponding ingress port of the core. Typically in the switch core itself the traffic cell is again queued in accordance with its destination switch port interface unit, and read out of the core queue at an appropriate time through the egress core port to the destination switch port interface unit.

For redundancy cell switching nodes typically operate with plural switch planes. Each of the plural switch planes has a switch core. Traffic cells are usually sent from the ingress switch port interface unit to the egress switch port interface unit through an active one of the plural switch planes. An example of an ATM cell switch with plural switch planes is disclosed in United States Patent Application SN 09/134,358 filed April 14, 1998 by Wicklund, entitled "Cell Selection for ATM Switch Having Redundant Switch Planes", which is incorporated herein by reference.

When the active switch plane develops a problem, traffic cells are instead routed through one of the redundant switch planes. Detection of an failure of the active switch plane can be problematic. Moreover, change over from the failed plane (which is to no longer serve as the active plane) to a redundant plane (which becomes the active plane) can be complex and liable to cell loss.

One conventional method of detecting failure of a switch plane is to permit the switch port interface units to communicate with one another via handshaking packets or the like. As long as successful bidirectional handshaking occurs, it is presumed that there is full connectivity through the active switch plane, and therefore that the switch plane is error free. But such handshaking technique unwisely assumes that there is always a switch port interface unit connected to a particular egress port of the switch core. Such assumption may not be correct, as the egress port of the core may be open for future use. Moreover, the switch port interface units may operate a different transmission rates, as can occur when switch port interface units at the same switching node are of differing grades or of differing generations. If the switch port interface units operate at different transmission rates in a handshaking scenario, care must be taken to adjust the times of the handshaking packets, lest there be a timeout or failure to receive a sent handshaking packet. The timing considerations to protect against such a

possibility introduce complex and awkward implementations of the handshaking procedures.

What is needed, therefore, and an object of the present invention, is an efficient and simple technique for detecting failure of a switch plane in a multi-plane cell switch.

5 An advantage of the present invention is a technique for (upon detection of switch plane failure) transferring all traffic cells to a redundant plane without losing, corrupting, or confusing the order of the traffic cells.

BRIEF SUMMARY OF THE INVENTION

10 A cell switch node has a first switch plane; a second switch plane; a sender switch port interface unit; and a receiver switch port interface unit. The sender switch port interface unit sends maintenance cells to the receiver switch port interface unit via the first switch plane and the second switch plane. The maintenance cells are applied to the first switch plane and the second switch plane in a predetermined sequence. The receiver switch port interface unit detects an erroneous switch plane when the
15 maintenance cells from the sender switch port interface unit do not arrive in the predetermined sequence. The predetermined sequence preferably comprises maintenance cells applied to the first switch plane and the second switch plane in alternation. The maintenance cell includes a plane indicator which informs the receiver switch port interface unit whether the maintenance cell traveled through the first switch
20 plane or the second switch plane

In one aspect of the invention, the sender switch port interface unit sends cycles of maintenance cells to the receiver switch port interface unit. Each cycle comprises plural sets of the predetermined sequence. At least a portion of the maintenance cell of one of the sets of the cycle is inverted with respect to a corresponding portion of the
25 maintenance cell of another of the sets of the cycle. Preferably each cycle comprises two sets of the predetermined sequence of maintenance cells, and the inverted/non-inverted portion of the maintenance cell in its payload.

When the switching node of the present invention has N number of receiver switch port interface units, the sender switch port interface unit applies maintenance
30 cells to the N number of receiver switch port interface units in a cycle which

comprises: (1) non-inverted maintenance cells sent via the first switch plane to each of the N number of receiver switch port interface units; (2) non-inverted maintenance cells sent via the second switch plane to each of the N number of receiver switch port interface units; (3) inverted maintenance cells sent via the first switch plane to each of the N number of receiver switch port interface units; and (4) inverted maintenance cells sent via the second switch plane to each of the N number of receiver switch port interface units.

In addition to detecting missing maintenance cells, the receiver switch port interface unit detects an erroneous switch plane by performing error checking with respect to contents of a received maintenance cell.

BRIEF DESCRIPTION OF THE DRAWINGS

The foregoing and other objects, features, and advantages of the invention will be apparent from the following more particular description of preferred embodiments as illustrated in the accompanying drawings in which reference characters refer to the same parts throughout the various views. The drawings are not necessarily to scale, emphasis instead being placed upon illustrating the principles of the invention.

Fig. 1 is a schematic view of portions of a multi-plane cell switching node according to an embodiment of the invention.

Fig. 2A is a diagrammatic view showing transmission of maintenance cells according to a mode of the present invention during normal operation of the node of Fig. 1.

Fig. 2B is a diagrammatic view showing transmission of maintenance cells according to a mode of the present invention upon occurrence of a failure of a switch plane of the node of Fig. 1.

Fig. 3 is a schematic view of portions of a multi-plane, multi-switch port interface unit cell switching node according to an embodiment of the invention.

Fig. 4A is a diagrammatic view of an example format for a MPR register included in a switch port interface unit of the present invention.

Fig. 4B(2)(1) is an example format for an initialization cell of the present invention.

5 Fig. 4B(2) is an example format for a maintenance cell of the present invention.

Fig. 5A and Fig. 5B are diagrammatic views showing transmission of cycles of maintenance cells according to one example mode of the present invention.

Fig. 6 is a diagrammatic view illustrating a physical connection and a pair of physical connections in a switching node of the present invention.

10 Fig. 7 is a diagrammatic view showing basic actions performed in a change of switch plane operation.

Fig. 7A is a diagrammatic view showing more detailed steps according to a particular implementation for the plane change operation of Fig. 7.

15 Fig. 8 is a schematic view of portions of a multi-plane, multi-switch port interface unit cell switching node according to an embodiment of the invention for illustrating the implementation of Fig. 7A.

Fig. 9 is a diagrammatic view of a set of registers and timers included in a switch port interface unit according to the invention.

20 Fig. 10 is a diagrammatic view showing certain events involved in a plane change operation of Fig. 7A.

Fig. 11 is a diagrammatic view showing basic modes of a switch port interface unit.

Fig. 12 is a flowchart showing basic steps involved in receiving maintenance cells at a receiver switch port interface unit in accordance with a mode of the invention.

Fig. 13 is a schematic view of portions of a multi-plane, multi-switch port interface unit cell switching node according to an embodiment of the invention illustrated in the context of a node of a telecommunications system.

DETAILED DESCRIPTION

In the following description, for purposes of explanation and not limitation, specific details are set forth such as particular architectures, interfaces, techniques, etc. in order to provide a thorough understanding of the present invention. However, it will be apparent to those skilled in the art that the present invention may be practiced in other embodiments that depart from these specific details. In other instances, detailed descriptions of well known devices, circuits, and methods are omitted so as not to obscure the description of the present invention with unnecessary detail.

Fig. 1 illustrates an example switching node 20 of an embodiment of the invention. The switching node 20 comprise any type of cell switch, such as (for example), an ATM cell switch or an Internet Protocol (IP) router. The switching node 20 includes a first switch plane 22 (having a first plane switch core 23); second switch plane 24 (having a second plane switch core 25); sender switch port interface unit (SPIC) 26_S; and receiver switch port interface unit (SPIC) 26_R. In general, the sender switch port interface unit (SPIC) 26 sends traffic cells through one of the switch planes 22, 24 to the receiver switch port interface unit (SPIC) 26. The particular switch plane through which traffic cells (e.g., cells with user data) are routed is considered to be the active switch plane. The other switch plane is referred to as the passive switch plane and, as described herein, provides redundancy in the event of failure of the active switch plane. All switch port interface units have the same switch plane defined as the active switch plane.

From a redundancy perspective, the switch port interface units (SPIC) 26 have various modes of operation as illustrated in Fig. 11. These modes, described in more detail subsequently, include the configuration mode 11-1; start-up_A mode 11-2A (for first switch plane 22); start-up_B mode 11-2B (for second switch plane 24); normal_A mode 11-3A (for first switch plane 22); normal_B mode 11-3B (for second switch plane 24), forced_A mode 11-4A (for first switch plane 22); and forced_B mode 11-4B (for second switch plane 22).

At the particular moment shown in Fig. 1, first switch plane 22 is considered to be an active switch plane, while second switch plane 24 is considered to be a passive switch plane. At any moment of operation, traffic cells flow only through whichever switch plane is currently the active plane (as depicted by arrows 30 through first switch plane 22 serving as the active switch plane). Any traffic cells received at the switch plane currently serving as the passive switch plane are discarded.

In addition to traffic cells, the switching node 20 of the present invention also utilizes the following cell types: initialization cells, management cells; maintenance cells; and plane change cells. Management cells (used, e.g., to control the switch cores 23, 25) can be directed to either of first the active switch plane or the passive switch plane. The maintenance cells are employed to detect a failed switch plane, and also can travel through both the active switch plane and the passive switch plane 24 as indicated by arrows 32 in Fig. 1. The plane change cells typically only travel through the passive switch plane, as shown by arrows 34 in Fig. 1.

The switching node 20 has four major error detection scenarios. The first error detection scenario occurs when the Access Multiplexer Switch Core Interface (ASCI) ingress goes out of synchronization, as detected by the Access Multiplexer Switch Core Circuit (ASCC) and reported to switch port interface unit (SPIC) 26. The second error detection scenario involves the ASCI going out of synchronization (as detected by switch port interface unit (SPIC) 26). A third error detection scenario is occurrence of a parity error in certain cells in the link from the ASCC to switch port interface unit (SPIC) 26 which are detected by the switch port interface unit (SPIC) 26 (e.g., the SPAS transfer list for traffic cells, in ASCU management cells, and switch maintenance cells). A fourth error detection scenario involves lost or corrupted switch maintenance cells. The fourth of the error detection scenarios, i.e., fault detection using maintenance cells, is discussed separately below.

If any of the four above-summarized error detection scenarios occurs in the active switch plane, then the switch port interface unit (SPIC) 26 which detects the error considers the plane to be erroneous. As explained in more detail below, the switch port interface unit (SPIC) 26 which detected the error will, when appropriate conditions are satisfied, immediately order every switch port interface unit (SPIC) 26 in the switching node 20 to change its designation of active switch plane. If any one of the above error

scenarios is detected in the passive switch plane, the error will be noted in a corresponding error register/counter, and the right to initiate a plane change will be withdrawn.

INITIALIZATION AND MAINTENANCE CELLS

A switch port interface unit (SPIC) 26 generates non-traffic cells internally that are sent to every switch port interface unit (SPIC) 26 (e.g., all sixty four switch ports in a sixty four switch port embodiment). The initialization cells are sent to reset the supervising mechanism in the receiving switch port interface unit (SPIC) 26 and to synchronize the switch ports cross-point status with the sending switch port interface unit's (SPIC) cross point information. The initialization cells are only sent during start-up. After start up there is an automatic change of initialization cells to maintenance cells. An example format of an initialization cell is illustrated in Fig. 4B(1).

Whereas synchronization error checking and parity error checking cover only errors in the interface between one switch port interface unit (SPIC) 26 and its corresponding part in the ASCC, switch maintenance cells are used to detect errors in the complete path between one switch port interface unit (SPIC) 26 and another switch port interface unit (SPIC) 26. An example format of switch maintenance cell is illustrated in Fig. 4B(2). As explained subsequently, to check the cross point buffers in the switch ports for stuck-at-zero or stuck-at-one faults, every other cell to a plane has its payload bit pattern inverted. The switch maintenance cells are also used to spread information about what planes a switch port interface unit (SPIC) 26 considers to be erroneous and which plane the switch port interface unit (SPIC) 26 considers as its active plane.

Concerning the format of the initialization cell (Fig. 4B(1)) and the format of the switch maintenance cell (Fig. 4B(2)), in the ingress direction the field "destination/source address" contains the destination address. In the egress direction, this field contains the source address. The logic address "111100" indicates that the cell is a switch maintenance cell or an initialization cell. The field "inv" reveals whether the payload pattern is inverted or not (e.g., 0= inverted; 1 = not inverted). The SMC/INI field are used to designate the cell as a switch maintenance cell or an initialization cell. That is, if the SMC/INI field has the value "0", the cell is a switch maintenance cell, but

- if the SMC/INI field has the value "1", the cell is an initialization cell. The field SPIC_ASP indicates whether the plane which the cell through which the cell is to be passed is considered an active plane or a passive plane (the field SPIC_ASP set to zero indicates an active plane; the field SPIC_ASP set to one indicates a passive plane). A bit set to one in the SPIC_plane_status bit indicates that the switch port interface unit (SPIC) 26 considers the corresponding switch plane as faulty.

FAULT DETECTION USING MAINTENANCE CELLS

- In switching node 20, each switch port interface unit 26_s serves as a sender switch port interface unit (SPIC) 26 to send switch maintenance cells to all corresponding cross points in first switch plane 22 and to all corresponding crosspoints in second switch plane 24. As explained subsequently, these switch maintenance cells are applied to the first switch plane 22 and the second switch plane 24 in a predetermined sequence to facilitate fault detection. The cells that have been sent to crosspoints that have switch ports connected to them will be delivered to the connected switch ports.

- In one mode of the invention, the predetermined sequence is an alternation of the maintenance cells. In this respect, in Fig. 2A the maintenance cells sent by sender switch port interface unit 26_s via first switch plane 22 are represented with the letter A while the maintenance cells sent by sender switch port interface unit 26_s via second switch plane 24 are represented with the letter B. Fig. 2A also shows that sender switch port interface unit 26_s sends the maintenance cells A, B in the predetermined sequence (e.g., alternation). Accordingly, for the normal operation scenario of Fig. 2A the receiver switch port interface unit 26_R receives the maintenance cells in the expected predetermined sequence (e.g., A,B,A, etc.).

- Fig. 2B shows a scenario in which first switch plane 22 fails, with the consequence that although sender switch port interface unit 26_s sends the maintenance cells in the same predetermined sequence as in Fig. 2A (A,B,A, etc.), the pattern of maintenance cells detected by receiver switch port interface unit 26_R is not the predetermined sequence. Rather, as shown in Fig. 2B, the receiver switch port interface unit 26_R detects the maintenance cells with the pattern A,B,B,A,B. In view of the fact that there is a break in the predetermined sequence of maintenance cells, receiver switch port interface unit 26_R concludes that the absence of a maintenance cell from first

switch plane 22 indicates a failure of first switch plane 22. Thus, receiver switch port interface unit 26_R detects an erroneous or failed switch plane when the maintenance cells from the sender switch port interface unit do not arrive in the predetermined sequence.

The foregoing figures and description have, for sake of simplicity, shown the switching node 20 as including only one sender switch port interface unit and one receiver switch port interface unit. As illustrated in Fig. 3, a typical switching node 320 has plural switch port interface units, such as sender switch port interface unit 26₁ through 26_N. In one example embodiment, the number of switch port interface units N is 64. Although each switch port interface unit functions dually as a sender switch port interface unit and a receiver switch port interface unit, the present discussion now concentrates on transmission of maintenance cells from just one switch port interface unit, in particular switch port interface unit 26₁. Thus, in the ensuing discussion, switch port interface unit 26₁ is considered to be the sender switch port interface unit 26_R, and the remaining switch port interface units (26₂ through 26_N) are considered to be the receiver switch port interface units relative to maintenance cells issued from sender switch port interface unit 26_R.

Since the sender switch port interface unit 26_R is, in the switching node 320 of Fig. 3, sending maintenance cells to plural receiver switch port interface units, each maintenance cell must have, e.g., an address to specify the particular receiver switch port interface unit to which it is directed. In the example embodiments of the present invention, the address and various other information to be included in the maintenance cell is prepared in a register maintained by the sender switch port interface unit. Fig. 4A shows example contents of such a register 40, known as the maintenance pointer register (MPR register). The MPR register 40 is used to keep track of what cell type to send and to which address in which switch plane the maintenance cell is to be sent. The MPR register 40 has several fields, including inversion status bit 42; plane indicator bit 44; and crosspoint address field 46. In the illustrated embodiment, the inversion status bit 42 is the most significant bit, the plane indicator bit 44 is the second most significant bit, and the crosspoint address field 46 occupies the six least significant bits of register 40.

It is the crosspoint address field 46 which essentially serves as a receiver switch port interface unit address, since the crosspoint designated thereby indicates a corresponding switch port interface unit. The MPR register 40 is operated as a wrap-around counter, in which crosspoint address field 46 is incremented each time a
 5 maintenance cell is issued from the sender switch port interface unit.

The plane indicator bit 44 advises whether the maintenance cell is to be sent through first switch plane 22 or second switch plane 24. For example, if the plane indicator bit 44 is set to "0", the maintenance cell is to be sent through first switch plane 22. Conversely, if plane indicator bit 44 is set to "1", the maintenance cell is to be sent
 10 through second switch plane 24. The inversion status bit 42 of MPR register 40 indicates whether the payload of the maintenance cell to be generated is to be inverted or non-inverted.

When a maintenance cell is issued from sender switch port interface unit 26, the information maintained in MPR register 40 at that moment is included in the
 15 maintenance cell. Reference is again made to Fig. 4B(2) for the example contents of a maintenance cell 50 as including a maintenance cell header 51; inversion status (inv) bit 52; maintenance/initialization indicator (SMC/INI) 53; active switch plane indicator bit 54; destination/source address field 56; logic address field 57; payload 58; and parity bits 59. The parity bits 59 are four bits, each of the four parity bits being computed
 20 over a corresponding octet (the parity value for octet 4 also being calculated over the complete fourth octet). The payload 58 comprises a crosspoint test pattern.

In accordance with one aspect of the present invention, the sender switch port interface unit 26_R sends a cycle of maintenance cells, the cycle comprising plural sets of the predetermined sequence. In keeping with the provision of the inversion status (IS)
 25 bit 52, at least a portion of the maintenance cell of one of the sets of the cycle is inverted with respect to a corresponding portion of the maintenance cell of another of the sets of the cycle. That is, the payload of the inverted maintenance cell is the binary complement of the payload of the non-inverted maintenance cell.

As one example of the cyclical aspect of the invention, each cycle issued
 30 comprises two sets of the predetermined sequence of maintenance cells. Such example is illustrated by Fig. 5A which, for sake of simplicity, again shows only two switch port

interface units, i.e., sender switch port interface unit 26_S and receiver switch port interface unit 26_R. Fig. 5A shows that, in normal operation, the maintenance cells are sent in cycles C, e.g., cycle C₁, cycle C₂, and so forth. Each cycle C has two sets S, particularly set S_{1,1} and set S_{1,2} for cycle C₁; set S_{2,1} and set S_{2,2} for cycle C₂; etc. Each set S includes the predetermined sequence of maintenance cells, e.g., for the illustrated embodiment, one maintenance cell from first switch plane 22 followed by one maintenance cell from second switch plane 24. In Fig. 5A, the maintenance cells from first switch plane 22 are illustrated as cells "A" when the cell is a non-inverted cell and as cells "a" when the cell is an inverted cell. In analogous manner, the maintenance cells from second switch plane 24 are illustrated as cells "B" when the cell is a non-inverted cell and as cells "b" when the cell is an inverted cell. As recalled from the foregoing, an "inverted cell" has the contents of its payload 57 inverted with respect to the payload 57 of a non-inverted cell. Thus, the maintenance cells issued in the following order from sender switch port interface unit 26_S to receiver switch port interface unit 26_R: Cycle C₁ (A,B,a,b); Cycle C₂ (A,B,a,b); and so forth.

At the receiving switch port interface unit, the maintenance cells are checked for bit errors. The payload of the maintenance cells contain a known bit pattern designed to test the cross-point buffer RAMs of the switch core. The receiving switch port interface units check this pattern. The use of the inverted and non-inverted payload enable detection of stuck-at-one and stuck-at-zero conditions in the switch core crosspoint buffers. The receiving switch port interface unit also checks the parity of the cell header.

When there are N switch port interface units, each set S_N of a cycle C_N includes both a non-inverted maintenance cell and an inverted maintenance cell destined for each switch port interface unit. In other words, as illustrated in Fig. 5B, each cycle C_N comprises two sets S_{N1} and S_{N2}. The set S_{N1} includes: (1) non-inverted maintenance cells sent via the first switch plane (plane A) to each of the N number of receiver switch port interface units; (2) non-inverted maintenance cells sent via the second switch plane (plane B) to each of the N number of receiver switch port interface units. The set S_{N2} includes: (3) inverted maintenance cells sent via the first switch plane (plane A) to each of the N number of receiver switch port interface units; and (4) inverted maintenance

cells sent via the second switch plane (plane B) to each of the N number of receiver switch port interface units.

Thus, in the cycle C of Fig. 5A and the cycle CN of Fig. 5B, it is seen that the predetermined sequence (e.g., the maintenance cells alternating between switch planes) is maintained. Accordingly, the receiver switch port interface unit 26_R is able to discern switch plane failure in the same manner as above described with reference to Fig. 2B, e.g., by detection of a maintenance cell dropped from the predetermined sequence by a faulty switch plane.

The foregoing describes cycles of maintenance cells issued by one sender switch port interface unit 26_S. In actuality, each switch port interface unit sends cycles of maintenance cells in precisely the same manner as above described to the other switch port interface units of the switching node.

In the illustrated embodiment having $N = 64$ switch port interface units, within a time period of 100 ms two complete cycles (each cycle having 256 maintenance cells) have been sent on each physical connection that connects the sender switch port interface unit 26_S to every other switch port interface unit. The nominal time between two consecutive cells is provided by a timer known as the maintenance cell timer (MCT timer).

The number of maintenance cells sent per cycle is reflected by Expression 1.
The nominal time between two consecutive maintenance cells is 100ms/256, or 0.39 ms.

Expression 1:

$$2 \text{ celltypes} * 2 \text{ switch planes} * 64 \text{ switch port interface units} = 256$$

The maintenance cells of the present invention are sent by sender switch port interface unit 26_S regardless of whether there is traffic on that physical connection or not. The sender switch port interface unit 26_S does not know whether there is a switch port interface unit connected at the core port specified by the crosspoint address, nor whether all N number of crosspoints actually exist. The maintenance cells of the present invention have a higher priority than any traffic cell selected for sending from the sender switch port interface unit to a crosspoint unit in the switch core (regardless of

Quality of Service [QoS] of the traffic cell). A timer, known as the maintenance wait timer (MWT timer) is started when there is a request to send a maintenance cell to an occupied crosspoint buffer. If the MWT timer has expired, and the crosspoint buffer has still not been released, then the maintenance cell will be sent anyway.

- 5 For the most pessimistic case, the time that a maintenance cell has to wait can be calculated using Expression 2:

Expression 2:

$$\text{TimeoutValue} = (\text{MaxBitsPerCell} \times \text{MaxNoOfCells}) / (\text{LowestBitRate})$$

- 10 The MaxNoOfCells is the maximum number of cells that this maintenance cell has to wait before it can be sent. The most pessimistic value is 2*the maximum number of switch ports. For 388.8 Mbit/s (which is the lowest bit rate) 8*(64+6) bits/cell (including the row buffer release octets) and (2*64+4) cells (four extra cells times added to allow for the time it can take to send a PlusPrio cell and receive a RowBufferRelease cell), the time is 191 μ s. This time is shorter than the time between
15 maintenance cells which means that no queue is needed for the maintenance cells.

Given a maximum error detection time, the capacity reduction due to the maintenance cell load in the switch can be calculated with Expression 3.

Expression 3:

$$\text{CapacityRed} = (\text{NoOfCellTypes} \times \text{NoOfSwitchPorts} \times \text{BitsPerMaintenanceCell}) / (\text{BitRate} \times \text{MaximumErrorDetectionTime})$$

20

With the worst case specification, i.e., 100 ms maximum error detection time, 2 types of cells (inverted and non inverted), 64 switch ports, 8*(64+6) bits per maintenance cell and 388.8 Mbit/s bit rate, the capacity reduction is 0.19%.

- 25 Thus, all physical connections through the first plane switch core 23 and the second plane switch core 25 are checked. As used herein, a physical connection is a unidirectional connection between a sender switch port interface unit and a receiver switch port interface unit. A pair of physical connections is two physical connections

that has the same sender switch port interface unit and receiver switch port interface unit, with each of the physical connections of the pair running through different switch planes. Fig. 6 illustrates both a physical connection and a pair of physical connections.

At each receiver switch port interface unit 26_R there is a maintenance process which checks the incoming maintenance cells. The maintenance process checks for both corrupted and lost maintenance cells. Corrupted maintenance cells are detected by performing error checking on the cell using, e.g., cyclical redundancy check field 58. If a corrupted maintenance cell is detected, the switch plane from which it came is considered erroneous. Lost maintenance cells are detected in the manner of Fig. 2A and Fig. 2B, e.g., by checking whether arrival of the maintenance cells is in the predetermined sequence (e.g., in an illustrated embodiment, alternating between the switch planes for each pair of physical connections). If the predetermined sequence is not maintained, the switch plane which omitted a maintenance cell from the sequence is considered to be erroneous. That is, the switch plane where the maintenance cell got lost is considered to have failed.

SENDING SPIC INFORMATION

The switch maintenance cells are not only used for error detection as above described, but also serve as carriers of information between switch port interface units (SPIC) 26. Three bits in the switch maintenance cell, illustrated in Fig. 4B(2), are used for this purpose. The SPIC_ASP bit is a copy of the ASP register in the switch port interface unit (SPIC) 26. The SPIC_plane_status_A and SPIC_plane_status_B bits reflect the opinion of the switch port interface unit (SPIC) 26 about first switch plane 22 and second switch plane 24. In this regard, The switch port interface unit (SPIC) 26 considers a switch plane as fully adequate if the following criteria are met: (1) the switch plane is configured; (2) the switch plane and the switch port have been in synchronization for a certain time; (3) no parity errors have occurred over a certain measured time in the SPAS transfer list; (4) no switch maintenance cells have been lost over a certain time; and (5) if a switch port interface unit (SPIC) 26 is forced to use a certain switch plane, it considers the passive switch plane as faulty.

When a switch port interface unit (SPIC) 26 detects errors according to criterion (2) and (3) in the preceding paragraph, switch port interface unit (SPIC) 26 provides a

corresponding indication in the switch maintenance cells. This indication remains as long as the error remains and during four additional full switch maintenance cell send cycles.

Every time a switch maintenance cell or initialization cell arrives at a switch port interface unit (SPIC) 26, the register bit of the corresponding sending switch port interface unit (SPIC) 26 is set in the AMC register. The switch port interface unit (SPIC) 26 also updates the plane_status and SM_ASP_status bits with the information found in the SPIC_plane_status and SPIC_ASP bits in the switch maintenance cell.

CHANGING ACTIVE SWITCH PLANE

One scenario for detecting a faulty switch plane has been described in detail above (the scenario using switch maintenance cells). When a faulty switch plane has been detected using any of the four above-summarized error detection scenarios (including the scenario using switch maintenance cells), the switching node 20 next executes a change of switch plane operation. Basic actions performed in the change of switch plane operation are depicted in Fig. 7, which is below discussed in the context of the switching node 20 of the embodiment of Fig. 3. It is assumed in the ensuing discussion of Fig. 7 that a fault has been detected in first switch plane 22 (which has been serving as the active switch plane), and that second switch plane 24 has been serving as the passive switch plane. The objective of the change of switch plane operation of Fig. 7 is to reverse or change roles of the first switch plane 22 and the second switch plane 24, so that the second switch plane 24 becomes the active switch plane and first switch plane 22 becomes the passive switch plane.

The change of switch plane operation of Fig. 7 begins with a fault detection in the active switch plane, as represented by arrow 7-1 in Fig. 7. When the passive switch plane realizes that a fault has occurred in the active switch plane (action 7-2), the passive switch plane prepares plane change cells (action 7-3) which are sent to each of the switch port interface units (SPIC) 26 as action 7-4 (only one of the switch port interface units 26 being shown in Fig. 7). Upon receipt of the plane change cell of action 7-4, as action 7-5 each switch port interface unit (SPIC) 26 sends a cell to suspend egress of cells from the second (passive) switch plane so that the passive switch plane enters a blocked egress state 7-7. The sending of such a egress-suspending cell is indicated as action 7-6 in Fig. 7.

In view of the fact that the (now faulty) active switch plane likely still has some traffic cells destined to the switch port interface unit 26, the switch port interface unit 26 continues to receive traffic cells from the active switch plane until the switch port interface unit has detected (as action 7-8) that all traffic cells destined to it have been
 5 flushed from the active switch plane. Upon detecting the complete flushing of the active switch plane, as action 7-9 the switch port interface unit (SPIC) 26 sends a cell which notifies the second switch plane that egress traffic flow from the second switch plane can resume. Thus, at action 7-10, the second switch plane 24 begins to output cells.

10 The switch port interface units, in the meanwhile, as action, 7-11 redirect their outgoing traffic cells to the passive switch plane rather than to the active switch plane as formerly. Since the second switch plane is now permitted to transmit traffic cells, these traffic cells emanating from switch port interface units (SPIC) 26 are routed through the second switch plane. Thus, the change of active switch plane is completed.

15 Fig. 7 thus shows general actions involved in the change of switch plane operation. One particular implementation of the change of switch plane operation is now described and illustrated in Fig. 7A in conjunction with Fig. 10, and in the context of the particular switching node 820 shown in Fig. 8.

As in the embodiment of Fig. 3, the switching node 820 has first switch plane 22,
 20 second switch plane 24, and plural switch port interface units $26_1 - 26_N$. In switching node 820 of Fig. 8 each switch plane has both a switch core port 80 and a cell receiving unit (RCU) 82 associated with each switch port interface unit 26. For example, with reference to switch plane 22, which is also shown as the "A" switch plane, the switch port interface unit 26_1 is connected by a bidirectional link to switch core port $80A_1$
 25 through which cells are communicated to cell receiving unit (RCU) $82A_1$ (the cell receiving unit (RCU) $82A_1$ being thusly associated with switch port interface unit 26_1). Similarly, with reference to switch plane 24, which is also shown as the "B" switch plane, the switch port interface unit 26_1 is connected by a bidirectional link to switch core port $80B_1$ through which cells are communicated to cell receiving unit (RCU) $82B_1$
 30 (the cell receiving unit (RCU) $82B_1$ being thusly associated with switch port interface unit 26_1). Moreover, each switch port interface unit 26 has a set 84 of registers and timers, with various registers and timers of the set 84 being utilized in one or both of the

switch plane fault detection and the change of switch plane operation. The contents of the set 84 of registers and timers is illustrated in Fig. 9, and described hereinafter.

In the implementation of Fig. 7A, it is assumed that initially first switch plane 22 is the active switch plane and that second switch plane 24 is the passive switch plane.

As described above in connection with the fault detection operation which utilizes a predetermined sequence of maintenance cells, in the Fig. 7A implementation it is one of the switch port interface units 26 which detects a fault in first (now active) switch plane 22, as indicated by action 7-1A. Upon such fault detection and provided that three conditions are satisfied, as action 7-1B the fault-detecting switch port interface unit 26 sends a fault detection cell to second (now passive) switch plane 24. The three conditions that must be satisfied before the sending of a fault detection cell are: (1) the second (passive) switch plane is in synchronization; (2) the passive plane has been configured; and (3) the error detecting receive unit is authorized to send a plane change cell (by the setting of a flag PC0_armed = true). The fault detection cell has a higher priority than the switch maintenance cells.

When second switch plane 24 recognizes (action 7-2) that a fault has occurred in first switch plane 22, as action 7-2A the second switch plane 24 broadcasts a plane change signal to each of the cell receiving units (RCUs) 82B of second switch plane 24. The plane change signal is carried on a signaling network 86B which connects all cell receiving units (RCUs) 82B of second switch plane 24.

Upon receipt of the plane change signal, each cell receiving unit (RCU) 82B of second switch plane 24 generates a plane change cell (action 7-3), and sends such plane change cell to its associated switch port interface unit 26 (action 7-4). When the plane change cell is received by the switch port interface unit 26 to which it is destined (see event 10-1 in Fig. 10), as action 7-5A the switch port interface unit (SPIC) 26 generates and sends to the second switch plane 24 a cell (known as the Egress_Traffic_Stop cell) which instructs the second switch plane 24 to stop egress of traffic flow from second switch plane 24 (see event 10-2 in Fig. 10). As action 7-5B, the switch port interface unit (SPIC) 26 starts the plane change timer (PCT) (see event 10-3 in Fig. 10). Further, as action 7-5C, switch port interface units (SPIC) 26 disable cell ingress from switch port interface unit (SPIC) 26 to both first switch plane 22 and second switch plane 24.

As long as the plane change timer (PCT) has not expired, ASCU ingress is disabled for both first switch plane 22 and second switch plane 24. This means that no traffic cells, switch maintenance cells, initialization cells, or plane change cells can be sent. The plane change timer (PCT) runs long enough to ensure that the slowest switch port interface units (SPIC) 26 have had time to receive the plane change cell and start their plane change process, e.g., stopped their ASCI ingress flows and sent an Egress_Traffic_Stop cell.

Three actions occur upon expiration of the plane change timer (PCT). Expiration of the plane change timer (PCT) is indicated as event 10-4 in Fig. 10. A first such action (action 7-5D) is that cell ingress is resumed. However, as further indicated by action 7-11, the traffic cell flow is now redirected to the former passive (now active) plane, i.e., second switch plane 24. As such, the switch maintenance cells should now have the new active plane indicated in their switch maintenance SPIC_ASP field.

As a second action (action 7-5E) occurring upon expiration of the plane change timer (PCT), a synchronization cell wait timer (SWT) is started (at event 10-4 in Fig. 10).

As a third action (action 7-8A) occurring upon expiration of the plane change timer (PCT), the switch port interface unit (SPIC) 26 starts to monitor the traffic out of the previous active plane (e.g., first switch plane 22). When (as action 7-8B) a series of a first predetermined type of non-traffic cells are received from the ASCC egress, the first switch plane 22 is considered to be empty. In the illustrated embodiment, action 7-8B involves detecting thirty two consecutive synchronization cells. Upon detection of the series of action 7-8B (see event 10-5 in Fig. 10), as action 7-8C the switch port interface unit (SPIC) 26 sends a cell (known as the Egress_Traffic_Start cell) to second switch plane 24 to restart the flow of cells from second switch plane 24. Action 7-9 and event 10-6 reflect transmission of the Egress_Traffic_Start cell to second switch plane 24. Upon receipt of the Egress_Traffic_Start cell, as action 7-10 egress of traffic flow is restarted from the second switch plane 24 (see event 10-7 in Fig. 10). In this way, the sequential correctness in the cell flow at the receiving switch port interface units (SPIC) 26 is guaranteed, and traffic cells do not arrive simultaneously from both first switch plane 22 and second switch plane 24.

Thus, the detection of flushing of all traffic cells destined to it from first switch plane 22 by switch port interface unit 26 involves actions 7-7A through 7-7C in the Fig. 7A implementation. Since the actions of Fig. 7A are being performed by all other switch port interface units (since they also received plane change cells), all other switch port interface units are also sending synchronization cells to all other switch port interface units. When, at action 7-8B, it is determined that synchronization cells have been received from all other switch port interface units, as action 7-8C the switch port interface unit realizes that the first switch plane 22 has been fully flushed of all traffic cells headed to the switch port interface unit (so that no traffic cells will be lost or left in first switch plane 22 during the change of switch plane operation).

From the foregoing, it should be apparent that the switch port interface units (SPIC) 26 that are fast in the completion of the plane change process, and which already have enabled their ASCII egress on the new active plane, can therefore start to receive traffic cells from other fast SPICs without having to wait for slow switch port interface units (SPIC) 26, thereby minimally disturbing traffic flow.

In the embodiment of Fig. 7A, the plane change timer (PCT) is set to have a predetermined time which is sufficiently long to ensure that the slowest switch port interface unit has had time to receive the plane change cell and to redirect traffic cells to the second switch plane.

All switch ports that receive a plane change cell will immediately lose their rights to send a plane change cell, i.e., the flag PCO_armed is set false to prevent oscillations between planes if there are multiple errors in the system.

During the plane change process and shortly thereafter, the switch port interface unit (SPIC) 26 will likely receive some switch maintenance cells with the ASP bit set to the old value. This will generate invalid SM status bits in the switch port interface units (SPIC) 26. This false indication of an error after plane change has to be taken into account.

If a plane change cell is received from an unconfigured plane, the plane change procedure is fully executed.

If, for some reason, a plane change cell arrives from the active switch plane rather than the passive switch plane, the switch port interface unit receiving the plane change cell will perform the change of switch plane operation, except that it will not change switch planes. In this situation, the switch port interface unit has to go through the change of switch plane operation since there might be other switch port interface units in the switch that actually do need to change plane.

PLANE CHANGE TIMER (PCT)

The purpose with the PCT is to make sure that no traffic is sent to a plane while the intended receiver consider the plane as passive.

The expiration value for the PCT timer is calculated as the sum of each of (1) - (7) below:

(1) The maximum time difference in the signalling when the switch port internally broadcasts to a plane change to all RCUs.

(2) The maximum time difference between ASCC ports when they finish the current send cycle.

(3) The maximum time difference when sending plane change cell.

(4) The maximum time difference between SPICs when they receive and recognize the plane change order.

(5) The maximum time difference between SPICs when they finish the current send cycle.

(6) The maximum time difference between SPICs when sending EgrTrafficStop cell.

(7) The maximum time difference between RCUs when stopping traffic in the egress port.

Assuming all cells are equally large and maximum differences internally in SPIC and ASCC is 0.5 cell send cycle gives:

$$0.5 + 1 + 0.5 + 0.5 + 2 + 0.5 + 0.5 = 5.5 \text{ cell cycles.}$$

With cell size 64 octets and a lowest bit rate of 388.8 Mbit/s that corresponds to
 5 $(5.5 * 64 * 8 / 388.8) \approx 7 \mu\text{s}$.

SYNC CELL WAIT TIMER (SWT)

The expiration value for the SWT timer is calculated as the sum of (1) the time it takes to empty the crosspoint buffer column in ASCC, and (2) the time it takes to send row buffer release cells between each traffic cell. Thus, for (64 traffic cells + 64
 10 rbr cells) = $(64 * 64 + 64 * 6)$ octets = 4480 octets, with a 16 bit data bus it will be 2240 clock cycles, independent of clock speed. SWT will have following values dependent of SAI clock speed.

If 48.6Mhz clock is used, $\text{SWT} = \text{Number of clock cycles} / \text{frequency}$
 $= 2240 / 48.6 \text{Mhz} = 47 \mu\text{s}$

15 If 24.3Mhz clock is used, $\text{SWT} = \text{Number of clock cycles} / \text{frequency}$
 $= 2240 / 24.3 \text{Mhz} = 93 \mu\text{s}$.

SWITCH PORT MODES

The switch port interface units 26 of the present invention basically function in the modes illustrated in Fig. 11. As mentioned previously, these modes include the
 20 configuration mode 11-1; start-up_A mode 11-2A (for first switch plane 22); start-up_B mode 11-2B (for second switch plane 24); normal_A mode 11-3A (for first switch plane 22); normal_B mode 11-3B (for second switch plane 24), forced_A mode 11-4A (for first switch plane 22); and forced_B mode 11-4B (for second switch plane 22).

25 The purpose for the configuration mode 11-1 is to give SPAS-SW time to perform necessary configurations of registers in the two switch planes and the SPIC before SPIC starts to send cells through the switch cores. The configuration mode 11-1

and after hardware reset. In configuration mode 11-1 a switch port interface unit will not send any traffic/switch maintenance/initialization or internally mgmt cells over ASCI ingress plane A or plane B. Only mgmt cells from SPAS-SW can be sent over ASCI ingress. Over ASCI egress every received cell, except for the ASCU_mgmt cells that are sent to the SPAS-SW, are discarded. Plane change cells should not start any plane change mechanism in the SPIC. The configuration of a switch port is finalized by setting the ASCC_config_attempt register. If no ASCI_clock_loss is indicated the corresponding switch plane, ASCC_configured bit is set to true. After SPAS-SW has finished the configuration of the registers, the SPAS-SW must send a SPIU_mgmt cell to set the SCM (SPIC_in_Config_Mode) register to the value config_ready. SPIC will then enter the start-up mode. It is possible for SPAS-SW to set the SCM register to the value config_mode at any time.

The start-up mode 11-2 is entered when start-up mode force_active_plane is set to redundancy_auto, then SCM is set to config_ready. There are five purposes for the start-up mode 11-2:

1. To decide active switch plane by checking the SPIC_ASP bit in the received maintenance cells. The SPIC changes from startup_A to startup_B if any received maintenance cell indicates that its sender has the B plane as active switch plane, and vice versa. If no switch maintenance cells are received, plane A will be considered as active plane.

2. To adapt to the switch maintenance cell sending order of the other switch port interface units (SPIC) 26 by setting the MCA register correctly.

3. To send initialization cells to make it possible for other switch port interface units (SPIC) 26 to set their MCA registers correctly before they start checking for lost maintenance cells from this SPIC.

4. To identify the Switch Module status (This means that the SPIC in start-up mode monitors the incoming maintenance cells from SPICs in normal mode. Two things are checked: (a) if any switch port interface unit (SPIC) 26 in normal mode indicates that it consider a switch plane as faulty, or (b) switch port interface units

(SPIC) 26 in normal mode are having different active switch plane. If any of these states are found the plane change order mechanism will be disabled (PCO_armed=false)).

- 5 5. To synchronize the row buffer status information in the switch port interface unit (SPIC) 26 with the actual status in the cross points in the switch ports.

In the start-up mode, only presync, sync, plusPrio, ASCU management and initialisation cells can be sent over ASCI ingress. No traffic cells, maintenance cells, or plane change cells will be sent.

- 10 Traffic cells received at ASCI egress are discarded, but embedded row buffer release information have to be used to update the row buffer status information in the SPIC. ASCU management cells are sent to SPAS-SW as in normal mode. Arriving switch maintenance cells and initialization cells are used only for updating the MCA register. Plane change orders are carried out.

- 15 Before putting itself into normal mode 11-3, the switch port sends at least two cycles of initialization cells, i.e. 512 cells. This is controlled by a cell counter called STNC, Start-up to Normal Counter. The counter is started first when the startup-mode is entered. The timer is restarted every time when a switch maintenance cell is received in which the SPIC_ASP bit does not indicate the same active plane as the previously received maintenance cell. The timer is also restarted if a plane change cell is received.
- 20 When the conditions are stable, the timer will eventually expire and the normal mode 11-3 will then be entered. In the special case where the timer never reaches its end value because it is repeatedly restarted, the normal mode 11-3 will never be entered.

- 25 The normal mode 11-3 occurs when the switch port is up and running and traffic cells can be transmitted and received. No traffic, switch maintenance, or initialization cells should be transmitted to unconfigured switch planes (although cells can be received). During plane change the traffic is turned off, as previously described (see Fig. 7 and Fig. 7A). During normal mode 11-3 the switch port continuously checks the status of the active plane. If the PCO_armed register is set (sending plane change order is allowed). the switch port will send a plane change order cell if any of the following
- 30 errors are detected on the active plane: (1) ASCC out of sync; (2) SPIC out of sync;

(3) parity errors in the received cells; (4) erroneous or missing redundancy switch maintenance cells. When the switch port receives a plane change cell in normal mode the plane change process is started, as previously described.

The switch port interface unit (SPIC) 26 can be forced to Forced Active Plane Mode 11-4 (e.g., Forced_A or Forced_B mode) in systems where the redundancy mechanism is not used. The switch port interface unit (SPIC) 26 can only be set to forced mode from configuration mode 11-1. This is done by setting the force_active_plane to forced_A or forced_B and after that leaving configuration mode 11-1. Then the switch port interface unit (SPIC) 26 will use the indicated switch plane as its active switch plane. In this mode switch port interface unit (SPIC) 26 will not send any plane change cells or react to any incoming plane change cells. This means that the switch plane redundancy function is turned off. The switch maintenance cells should be sent as in normal mode with the switch plane not active marked as faulty. The ASP register will show the plane to which the SPIC is forced as active plane.

In both start-up mode 11-2 and normal mode 11-3 the switch port interface unit (SPIC) 26 has to monitor the incoming switch maintenance cells. Information in these cells are used for updating the AMC and the SM status register. If the incoming switch maintenance cells indicate that SPICs have different planes as active or that one or more SPICs have problem with at least one switch plane, the plane change order mechanism should be deactivated immediately, i.e. PCO_armed=false.

If the switch port interface unit (SPIC) 26 detects an ASCI_clock_loss the SPIC should set the corresponding ASCC_configured=false. This will give the SPAS-SW a mechanism for detecting new switch planes in the system. The SPAS-SW will then configure the ASCC port and acknowledge this to the SPIC.

SWITCH PORT INTERFACE UNITS: STATUS REGISTERS AND TIMERS

The timers and registers in set 84 of switch port interface unit (SPIC) 26, illustrated in Fig. 9, are discussed briefly below.

Plane Change Order Armed and Control register 9-1: The SPAS-SW can order the switch port interface unit (SPIC) 26 to activate the plane change mechanism by writing "activate" ('1') to the Plane Change Order Control, PCOC, bit. If the switch port

interface unit (SPIC) 26 activates the plane change order mechanism it sets the Plane Change Status, PCO_armed, bit to "true" ('1'). Deactivation is done by writing "deactivate" ('0') to the PCOC register. The plane change order mechanism can only be activated by SPAS-SW if the SPIC's SM status bits indicates that all SPICs can run traffic through any plane and that all planes have the same active plane. If any of these conditions should not be fulfilled at any moment, if a plane change cell is detected, or the SPIC sends a plane change cell the plane change order mechanism should be deactivated, i.e. PCO_armed= "false" ('0'). The plane change mechanism can only be reactivated by SPAS-SW by writing "activate" in the PCOC register. Note that it is not the value '1' in the PCOC register that activates the plane change mechanism, it is the write action combined with the necessary conditions for the setting of the PCO_armed "true" as stated previously. The PCOC is a control bit. The SW can only write to the register. The PCO_armed is a status bit and should be reset to false. The status bit can only be read by software.

ASCC_configured registers 9-2 and ASCC_config_attempt register 9-3:
Together these registers have four bits totally, including two status bits ASCC_configured_A and ASCC_configured_B and two control bits ASCC_config_attempt_A and ASCC_config_attempt_B. The switch port interface unit (SPIC) 26 has to know when an ASCC is configured and ready to receive cells without risking loss of cells. An ASCC is considered to be ready to receive cells when its corresponding ASCC_configured=true ('1'). The SW confirms that the ASCC is configured to the SPIC by writing to the related ASCC_config_attempt register. If the related ASCI_clock_loss status register indicates that the clock is present, the ASCC_configured register is set to "true". If an error occur which sets the ASCI_clock_loss, the SPIC will set the ASCC_configured register to "false". There are three states a SPIC can consider an ASCC to be in: No clock, clock and config. If the SPIC has an ASCI clock and the SW writes to ASCC_config_attempt, the SPIC consider the ASCC as configured. It is not the value '1' in the ASCC_config_attempt register that sets the ASCC_configured bits, it is the write action combined with the necessary conditions for the setting of the ASCC_configured "true" as stated previously. ASCC_config_attempt='0' has no meaning and should not affect any status bit. ASCC_config_attempt_A and ASCC_config_attempt_B are both control registers that are only writable. ASCC_configured_A and ASCC_configured_B are both readable

only status bits and should be reset to false.

Switch Module ASP register 9-4: The SM bits are compiled from corresponding bits in the switch maintenance cells. If the maintenance cell's SPIC_ASP indicates that the transmitting cell has ASP=A, then SM_ASP_status_A is set ('1'). If the maintenance cell's SPIC_ASP indicates that the transmitting cell has ASP=B, then SM_ASP_status_B is set ('1'). Note that the values in the SM_ASP is latched, i.e. the register value contains information measured from the last reset. This might lead to false error indications during plane change, for instance. After SW has reset the register, the register will not be updated correctly until after 25ms. The registers are read and reset to '0' by SPAS-SW. Reset value is '0'.

Plane status register 9-5: The plane status register 9-5 is actually two 64 bit registers, namely Plane_status_A and Plane_status_B. These registers are updated with information extracted from the maintenance cells SPIC_plane_status bits. A high level in a register position indicates that at least one switch maintenance cell with a '1' in the plane status field has been received from the corresponding SPIC. The values in the plane_status register are latched, i.e. the register value contains information measured from the last reset. This might lead to false error indications during plane change, for instance. After SW has reset the register, the register will not be updated correctly until after 25ms. The registers are read and reset to '0' by SPAS-SW. The reset value is '0'.

Push to normal mode register 9-6: The normal mode register 9-6 is a one bit register. Writing '1' to the push_to_normal_mode register forces the SPIC to Normal_A mode. This should only be possible from startup mode. The register has no meaning when the SPIC is in normal mode or forced. A transition from startup to normal mode caused by SW push should not lead to any error indications in maintenance cell supervising mechanisms. This register is only writable. Writing '0' is not considered as valid.

Active Switch Plane (ASP) register 9-7: The Active Switch Plane (ASP) register 9-7 is a one bit register, specifying what switch plane is active. The register is set to "plane_A" ('0') after reset. The register can only be read, not written to or reset, by SPAS-SW. The switch port uses this register to determine in what switch plane it shall send and receive traffic cells. During plane change the SPIC will send and receive

traffic from different planes. The value of the ASP register is therefore not defined during plane change.

Forced Active Plane register 9-8: The Forced Active Plane register 9-8 is a two bit register used by the SPAS-SW to decide what mode the SPIC will run in after leaving config mode. The register is set to "redundancy_auto" ('00') after reset which will take the SPIC to startup mode when ordered to leave config mode. The register can be read and written to by SPAS-SW but should not influence the SPIC when the SPIC is in startup or normal mode. If the register is set to "forced_A" ('10') or "forced_B" ('11'), then the SPIC will be forced to plane A or B when leaving config mode. When the SPIC is forced to the A or B plane the register value '0x' is an invalid value that shall be considered the same as no change in state.

SPIC in Start-up Mode (SSM) register 9-9: The SPIC in Start-up Mode (SSM) register 9-9 is a one bit register showing whether SPIC is in normal or forced mode ('0'), or start-up or config mode ('1'). The register can be read, but not set/reset, by SPAS-SW. If SPAS-SW sets SCM to config_mode this register will indirectly be set to start-up mode 11-2 by SPIC.

SPIC in Config Mode (SCM) register 9-10: The SPIC in Config Mode (SCM) register 9-10 is a one bit register. The register is set to "config_mode" ('1') after reset. The value is '0' in all other modes. The register can be read and written to by SPAS-SW.

ASCI clock loss register 9-11: The ASCII clock loss register 9-11 is a two bit register. The register indicates a loss of clock in the ASCII egress interface. A high value indicates that the SPIC has lost the clock at least once since reset. The register can only be read and reset by SPAS-SW.

Arrived switch Maintenance Cells (AMC) 9-12: The Arrived switch Maintenance Cells (AMC) 9-12 is a 64 bit register. A bit is set when a switch maintenance cell or an initialisation cell has arrived from the corresponding switch port at any switch plane. All bits in the registers are set to zero after reset. By reading and resetting this register, the SPAS-SW can perform measurements to check which SPICs that are connected. SPAS-SW can read and reset the register.

Plane Change Ordered (PCO) register 9-13: The Plane Change Ordered (PCO) register 9-13 is a one bit register. If the register is set to "pcq_sent" ('1') it indicates that this switch port has initiated a change of active switch plane (has sent a plane change cell). The register can be read and reset by SPAS-SW.

5 Lost or corrupt switch maintenance cell (LMA and LMB) registers 9-14: The lost or corrupt switch maintenance cell (LMA and LMB) registers 9-14 are two 64 bits registers. Example of usage: if SPIC considers a switch maintenance cell to be corrupt or missing from plane B, sent from SPIC x, bit x will be set in the LMB register. The register can be read and reset by SPAS-SW.

10 The redundancy related SPIU timers 9-15 through 9-22 listed below are required for the implementation of the redundancy handling. The timers themselves cannot be accessed (read, written to or reset) by SPAS-SW. The expiration time for the timers can however be configured by SPIU_mgmt cells from SPAS-SW.

15 Maintenance Cell Timer (MCT) 9-15: The Maintenance Cell Timer (MCT) 9-15 expires when it is time to send the next switch maintenance cell or initialisation cell. The timer is restarted as soon as it has expired. The expiration value can be configured with the Maintenance Cell Timeout Register 9-19.

20 Maintenance Wait Timer (MWT) 9-16: If a crosspoint is occupied when it is time to send a switch maintenance cell or initialisation cell, the Maintenance Wait Timer (MWT) 9-16 is started. When it expires the cell will be sent even if the crosspoint still is occupied. The expiration value can be configured with the Maintenance Wait Timeout Register 9-20.

25 Plane Change Timer (PCT) 9-17: The Plane Change Timer (PCT) 9-17 is used during the plane change process. When this timer is running no cells, except sync and presync cells, may be sent in ASCI ingress in any plane. The expiration value can be configured with the Plane Change Timeout Register 9-21.

 Sync cell Wait Timer (SWT) 9-18: The Sync cell Wait Timer (SWT) 9-18 is used during the plane change process. The timer indicates for how long time a SPIC shall wait until all the crosspoint column buffers in the former active ASCC with

certainty have been unloaded. The expiration value can be configured with the Sync cell Wait Timeout Register 9-22.

Maintenance Cell Timeout Register 9-19: The value of the Maintenance Cell Timeout Register 9-19 decides when the Maintenance Cell Timer 9-15 shall expire, and therefore when it is time to send a new switch maintenance/initialisation cell. The register can be read and written by SPAS-SW. After reset the register shall have a default value. The resolution in the register is 1 μ s.

Maintenance Wait Timeout Register 9-20: The value of the Maintenance Wait Timeout Register 9-20 decides when the switch maintenance Wait Timer 9-16 shall expire, and therefore for how long SPIC shall wait before sending a switch maintenance/initialisation cell to an occupied crosspoint buffer. The register can be read and written by SPAS-SW. After reset the register shall have a default value. The resolution in the register is 1 μ s.

Plane Change Timeout Register 9-21: The value of the Plane Change Timeout Register 9-21 decides when the Plane Change Timer 9-17 shall expire. The register can be read and written by SPAS-SW. After reset the register shall have a default value. The resolution in the register is 1 μ s.

Sync cell Wait Timeout Register 9-22: The value of the Sync cell Wait Timeout Register 9-22 register decides when the Sync cell Wait Timer 9-18 shall expire. The register can be read and written by SPAS-SW. After reset the register shall have a default value. The resolution in the register is 1 μ s.

Maintenance Pointer Register (MPR) 9-23: The Maintenance Pointer Register (MPR) 9-23 is an eight bit register which indicates which type of maintenance cell shall be generated. A zero indicates a maintenance cell with a non-inverted payload bit pattern; a one indicates a maintenance cell with inverted payload bit pattern. The most significant -1 bit indicates a switch plane. A zero MSB indicates plane A; a one indicates plane B. The six least significant bits indicate a destination switch port (0-63). The MPR is operated as an incrementing wrap-around counter. The register is incremented each time a maintenance cell is sent.

Maintenance Cell Arrival Register (MCA) 9-24: The Maintenance Cell Arrival Register (MCA) 9-24 is a 64 bit register which represents a pair of physical connections through the switch. A bit is set to zero each time a maintenance cell arrives from switch plane A over the corresponding pair of physical connections. A bit is set to one each time a maintenance cell or an initialization cell arrives from switch plane B over the corresponding pair of physical connections. The MCA is used to determine whether a maintenance cell has been lost or not. Suppose, for example, that a maintenance cell arrives from the same switch plane as the preceding maintenance cell over the same pair of physical connections. In such case the maintenance cell has been lost since the maintenance cells are sent in alternating switch planes for a certain pair of physical connections. The Maintenance Cell Arrival Register (MCA) 9-24 only generates errors in the Forced Active Plane Mode 11-4 or the normal mode 11-3.

Start to Normal Counter (STNC) 9-25: The Start to Normal Counter (STNC) 9-25 is used to determine the transition between start-up mode 11-2 and normal mode 11-3. When this counter expires the switch port interface unit (SPIC) 26 will change from the start-up mode 11-2 to the normal mode 11-3. It is reset at plane change orders or when the switch port interface unit (SPIC) 26 changes from startup_A to startup_B mode, or vice versa.

SPIC CELL HANDLING

Fig. 12 shows example basic steps involved in a maintenance cell-handling procedure performed by a switch port interface unit. After the maintenance cell-handling procedure has begun (step 12-1), the arriving cell is checked at step 12-2 to determine whether it is one of a maintenance cell or an initialization cell. Other types of cells are ignored by the maintenance cell-handling procedure. Assuming that the arriving cell is one of a maintenance cell or an initialization cell, at step 12-3 the maintenance cell-handling procedure determines whether the arriving cell is a maintenance cell. If so, the maintenance cell-handling procedure checks the order of arrival of the arriving maintenance cell at step 12-4 by checking the Maintenance Cell Arrival Register (MCA). In this regard, the MCA is used to determine whether a maintenance cell has been lost or not. Suppose, for example, that a maintenance cell arrives from the same switch plane as the preceding maintenance cell over the same pair of physical connections. In such case the maintenance cell has been lost since the

maintenance cells are sent in alternating switch planes for a certain pair of physical connections.

If the arriving maintenance cell is not in the correct order of the predetermined sequence (see Fig. 2, for example), at step 12-5 the maintenance cell-handling procedure updates one of the Erroneous Cell Counters (ECA and ECB), as appropriate, depending from which switch plane the maintenance cell arrived. The Erroneous Cell Counters (ECA and ECB) count discarded cells, erroneous maintenance cells, and lost maintenance cells.

Upon an affirmative determination of step 12-4, at step 12-5 the maintenance cell-handling procedure determines whether the arriving maintenance cell has the correct cell contents, e.g., whether the arriving maintenance cell has proper integrity based on its parity values. If not, step 12-5 is performed for updating either the ECA or ECB counter.

After an affirmative determination of step 12-6, or after updating the ECA or ECB counter of step 12-5, at step 12-7 the Arrive Maintenance Cells (AMC) is updated. As explained previously the AMC is a 64 bit register which is set when a maintenance cell or an initialization cell has arrived from the bit-corresponding switch port interface unit. Then, as step 12-8, the maintenance cell-handling procedure updates the Maintenance Cell Arrival Register (MCA). As mentioned before, the MCA is a 64 bit register which represents a pair of physical connections through the switch. A bit is set to zero each time a maintenance cell arrives from switch plane A over the corresponding pair of physical connections. A bit is set to one each time a maintenance cell or an initialization cell arrives from switch plane B over the corresponding pair of physical connections. Then, at step 12-9, a sixty four bit register containing information on which ports are generating initialization or maintenance cells towards the receiving port is updated. by reading this register it is thus possible from the receiving port to see which other ports are present in the system.

SWITCH CORE

Details of one example implementation of switch core with its cell receiving units (RCUs) utilized for the present invention, including the first plane switch core 23

and the second plane switch core 25, are understood with reference to United States Patent Application Serial Number 09/188,101, entitled "Asynchronous Transfer Mode Switch", and United States Patent Application Serial Number 09/188,265, entitled "Asynchronous Transfer Mode Switch", both of which are incorporated herein by reference.

UTILIZATION IN TELECOMMUNICATIONS NODE

Fig. 13 shows an implementation of the present invention in a node 1320 having first switch plane 1322; second switch plane 1324; and switch port interface units 26_1 - 26_N . Each of the switch port interface units 26 is mounted on a device board 1370, device boards 1370_1 through 1370_N being shown for the mounting of respective switch port interface units 26_1 - 26_N . On the side of each switch port interface unit opposite the switch core, each switch port interface unit is connected by a bus or the like to one or more devices. In the illustration of Fig. 13, each device board 1370 has a respective board processor BP as one of its devices, and another device 1372. The devices 1372 can be, for example, exchange terminals handling a link that connects the node to the remainder of a network, e.g., to another cell switch node. The implementation of Fig. 13 can be utilized for a node of a telecommunications system 18, such as a base station node or a radio network controller (RNC) node [also known as a base station controller (BSC) node]. Examples of usage of such device boards are found in the following United States Patent applications (all of which are incorporated herein by reference): United States Patent Application Serial Number 09/188,102, entitled "Asynchronous Transfer Mode System Handling Differing AAL Protocols"; United States Patent Application Serial Number 09/188,097, entitled "Centralized Queuing for ATM Node"; United States Patent Application Serial Number 09/188,340, entitled "Cell Handling Unit and Method for ATM Node"; United States Patent Application Serial Number 09/188,347, entitled "ATM Time-Stamped Queuing"; United States Patent Application Serial Number 09/188,344, entitled "Coordinated Cell Discharge From ATM Queue"; and, United States Patent Application Serial Number 09/188,096, entitled "Combined Header Parameter Table for ATM Node". Preferably, but not exclusively, the cells utilized by present invention are ATM cells.

While the invention has been described in connection with what is presently considered to be the most practical and preferred embodiment, it is to be understood that the invention is not to be limited to the disclosed embodiment, but on the contrary, is intended to cover various modifications and equivalent arrangements included within
5 the spirit and scope of the appended claims.

FIG. 1
FIG. 2
FIG. 3
FIG. 4
FIG. 5
FIG. 6
FIG. 7
FIG. 8
FIG. 9
FIG. 10
FIG. 11
FIG. 12
FIG. 13
FIG. 14
FIG. 15
FIG. 16
FIG. 17
FIG. 18
FIG. 19
FIG. 20
FIG. 21
FIG. 22
FIG. 23
FIG. 24
FIG. 25
FIG. 26
FIG. 27
FIG. 28
FIG. 29
FIG. 30
FIG. 31
FIG. 32
FIG. 33
FIG. 34
FIG. 35
FIG. 36
FIG. 37
FIG. 38
FIG. 39
FIG. 40
FIG. 41
FIG. 42
FIG. 43
FIG. 44
FIG. 45
FIG. 46
FIG. 47
FIG. 48
FIG. 49
FIG. 50
FIG. 51
FIG. 52
FIG. 53
FIG. 54
FIG. 55
FIG. 56
FIG. 57
FIG. 58
FIG. 59
FIG. 60
FIG. 61
FIG. 62
FIG. 63
FIG. 64
FIG. 65
FIG. 66
FIG. 67
FIG. 68
FIG. 69
FIG. 70
FIG. 71
FIG. 72
FIG. 73
FIG. 74
FIG. 75
FIG. 76
FIG. 77
FIG. 78
FIG. 79
FIG. 80
FIG. 81
FIG. 82
FIG. 83
FIG. 84
FIG. 85
FIG. 86
FIG. 87
FIG. 88
FIG. 89
FIG. 90
FIG. 91
FIG. 92
FIG. 93
FIG. 94
FIG. 95
FIG. 96
FIG. 97
FIG. 98
FIG. 99
FIG. 100

WHAT IS CLAIMED IS:

1 1. A cell switch comprising:
2 a first switch plane comprising a first switch core;
3 a second switch plane comprising a second switch core;
4 a sender switch port interface unit;
5 a receiver switch port interface unit;
6 wherein the sender switch port interface unit sends maintenance cells to the
7 receiver switch port interface unit via the first switch plane and the second switch plane,
8 the maintenance cells being applied to the first switch plane and the second switch plane
9 in a predetermined sequence;
10 wherein the receiver switch port interface unit detects an erroneous switch plane
11 when the maintenance cells from the sender switch port interface unit do not arrive in
12 the predetermined sequence.

1 2. The apparatus of claim 1, wherein the predetermined sequence comprises
2 maintenance cells applied to the first switch plane and the second switch plane in
3 alternation.

1 3. The apparatus of claim 1, wherein the maintenance cell includes a plane
2 indicator which informs the receiver switch port interface unit whether the maintenance
3 cell traveled through the first switch plane or the second switch plane.

1 4. The apparatus of claim 1, wherein the sender switch port interface unit sends
2 a cycle of maintenance cells to the receiver switch port interface unit, the cycle
3 comprising plural sets of the predetermined sequence, and wherein at least a portion of
4 the maintenance cell of one of the sets of the cycle is inverted with respect to a
5 corresponding portion of the maintenance cell of another of the sets of the cycle.

1 5. The apparatus of claim 4, wherein the cycle comprises two sets of the
2 predetermined sequence of maintenance cells.

1 6. The apparatus of claim 1, wherein the receiver switch port interface unit
2 detects an erroneous switch plane by performing error checking with respect to contents
3 of a received maintenance cell.

1 7. The apparatus of claim 1, further comprising N number of receiver switch
2 port interface units, and wherein the sender switch port interface unit applies
3 maintenance cells to the N number of receiver switch port interface units in a cycle, the
4 cycle comprising (1) non-inverted maintenance cells sent via the first switch plane to
5 each of the N number of receiver switch port interface units; (2) non-inverted
6 maintenance cells sent via the second switch plane to each of the N number of receiver
7 switch port interface units; (3) inverted maintenance cells sent via the first switch
8 plane to each of the N number of receiver switch port interface units; and (4) inverted
9 maintenance cells sent via the second switch plane to each of the N number of receiver
10 switch port interface units, the inverted maintenance cells having at least a portion
11 thereof inverted with respect to the corresponding non-inverted cell of the cycle.

1 8. A cell switch comprising:
2 a switch plane comprising a switch core;
3 a sender switch port interface unit;
4 a receiver switch port interface unit;
5 wherein the sender switch port interface unit sends a pair of maintenance cells to
6 the receiver switch port interface unit via the switch plane, a first maintenance cell of
7 the pair having a predetermined portion thereof inverted with respect to a corresponding
8 portion of a second maintenance cell of the pair.

1 9. The apparatus of claim 8, further comprising a second switch plane through
2 which the sender switch port interface unit also sends a second pair of maintenance
3 cells to the receiver switch port interface unit, a first maintenance cell of the second pair
4 having a predetermined portion thereof inverted with respect to a corresponding portion
5 of a second maintenance cell of the second pair.

1 10. A cell switch comprising:
2 a first switch plane comprising a first switch core;
3 a second switch plane comprising a second switch core;
4 a sender switch port interface unit;
5 plural receiver switch port interface units, and wherein the sender switch port
6 interface unit applies maintenance cells to the plural receiver switch port interface units
7 in a cycle, the cycle comprising (1) non-inverted maintenance cells sent via the first
8 switch plane to each of the plural receiver switch port interface units; (2) non-inverted
9 maintenance cells sent via the second switch plane to each of the plural receiver switch
10 port interface units; (3) inverted maintenance cells sent via the first switch plane to
11 each of the plural receiver switch port interface units; and (4) inverted maintenance
12 cells sent via the second switch plane to each of the plural receiver switch port interface
13 units, the inverted maintenance cells having at least a portion thereof inverted with
14 respect to the corresponding non-inverted cell of the cycle.

1 11. A method of operating a cell switch comprising:
2 a first switch plane comprising a first switch core;
3 a second switch plane comprising a second switch core;
4 a sender switch port interface unit;
5 a receiver switch port interface unit;
6 a sender switch port interface unit sending maintenance cells to a receiver switch
7 port interface unit via a first switch plane and a second switch plane, the maintenance
8 cells being applied to the first switch plane and the second switch plane in a
9 predetermined sequence;
10 detecting, at the receiver switch port interface unit, an erroneous switch plane
11 when the maintenance cells from the sender switch port interface unit do not arrive in
12 the predetermined sequence.

1 12. The method of claim 11, further comprising forming the predetermined
2 sequence by alternating application of the maintenance cells to the first switch plane
3 and to the second switch plane.

1 13. The method of claim 11, further comprising including in the maintenance
2 cell a plane indicator and using the plane indicator at the receiver switch port interface
3 unit to determine whether the maintenance cell traveled through the first switch plane or
4 the second switch plane.

1 14. The method of claim 11, further comprising sending a cycle of maintenance
2 cells to the receiver switch port interface unit, the cycle comprising plural sets of the
3 predetermined sequence, and wherein at least a portion of the maintenance cell of one
4 of the sets of the cycle is inverted with respect to a corresponding portion of the
5 maintenance cell of another of the sets of the cycle.

1 15. The method of claim 14, further comprising forming the cycle with two sets
2 of the predetermined sequence of maintenance cells.

1 16. The method of claim 11, further comprising detecting at the receiver switch
2 port interface unit an erroneous switch plane by performing error checking with respect
3 to contents of a received maintenance cell.

1 17. The method of claim 11, further comprising the sender switch port interface
2 unit applying maintenance cells to N number of receiver switch port interface units in a
3 cycle, the cycle comprising (1) non-inverted maintenance cells sent via the first switch
4 plane to each of the N number of receiver switch port interface units; (2) non-inverted
5 maintenance cells sent via the second switch plane to each of the N number of receiver
6 switch port interface units; (3) inverted maintenance cells sent via the first switch
7 plane to each of the N number of receiver switch port interface units; and (4) inverted
8 maintenance cells sent via the second switch plane to each of the N number of receiver
9 switch port interface units, the inverted maintenance cells having at least a portion
10 thereof inverted with respect to the corresponding non-inverted cell of the cycle.


1 18. A method of operating a cell switch comprising sending a pair of
2 maintenance cells from a sender switch port interface unit to a receiver switch port
3 interface unit via a switch plane, a first maintenance cell of the pair having a
4 predetermined portion thereof inverted with respect to a corresponding portion of a
5 second maintenance cell of the pair.

1 19. The method of claim 18, further comprising sending a second pair of
2 maintenance cells via a second switch plane to the receiver switch port interface unit, a
3 first maintenance cell of the second pair having a predetermined portion thereof
4 inverted with respect to a corresponding portion of a second maintenance cell of the
5 second pair.

1 20. A method of operating a cell switch comprising applying maintenance cells
2 from a sender switch port interface unit to plural receiver switch port interface units in a
3 cycle, the cycle comprising (1) non-inverted maintenance cells sent via a first switch
4 plane to each of the plural receiver switch port interface units; (2) non-inverted
5 maintenance cells sent via a second switch plane to each of the plural receiver switch
6 port interface units; (3) inverted maintenance cells sent via the first switch plane to
7 each of the plural receiver switch port interface units; and (4) inverted maintenance
8 cells sent via the second switch plane to each of the plural receiver switch port interface
9 units, the inverted maintenance cells having at least a portion thereof inverted with
10 respect to the corresponding non-inverted cell of the cycle.

ABSTRACT OF THE DISCLOSURE

2 A cell switch node (20) has a first switch plane (22); a second switch plane (24);
 3 a sender switch port interface unit (26_S); and a receiver switch port interface unit (26_R).
 4 The sender switch port interface unit sends maintenance cells to the receiver switch port
 5 interface unit via the first switch plane and the second switch plane. The maintenance
 6 cells are applied to the first switch plane and the second switch plane in a predetermined
 7 sequence. The receiver switch port interface unit detects an erroneous switch plane
 8 when the maintenance cells from the sender switch port interface unit do not arrive in
 9 the predetermined sequence.



320






Fig 3



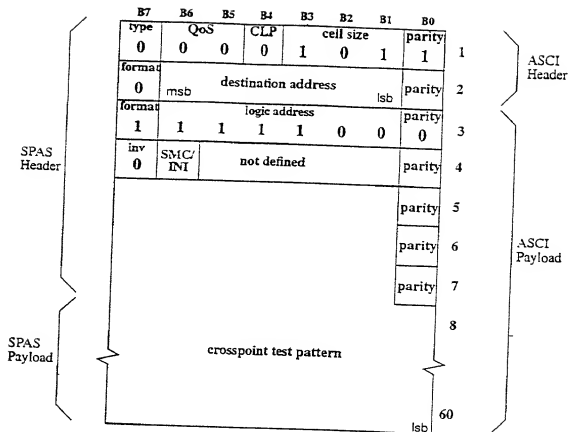


Fig 4B(1) Initialisation Cell

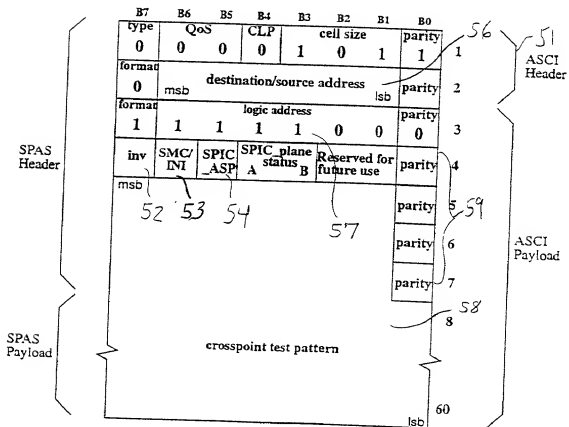


Fig 4B(2) Switch Maintenance Cell




Fig 5A

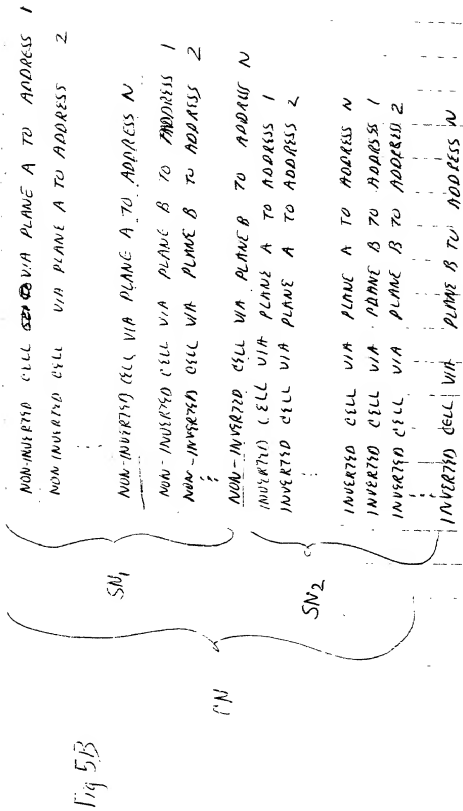


Fig 5B

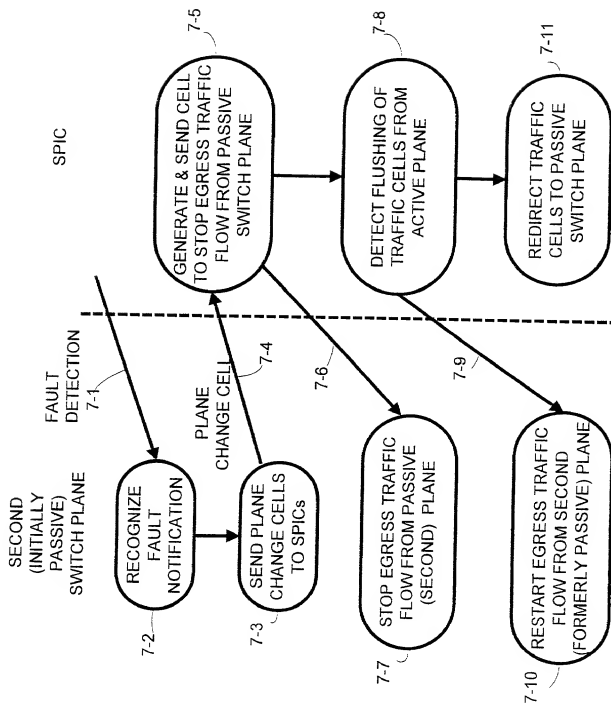


Fig. 7




Fig. 7A

820




Fig. 8

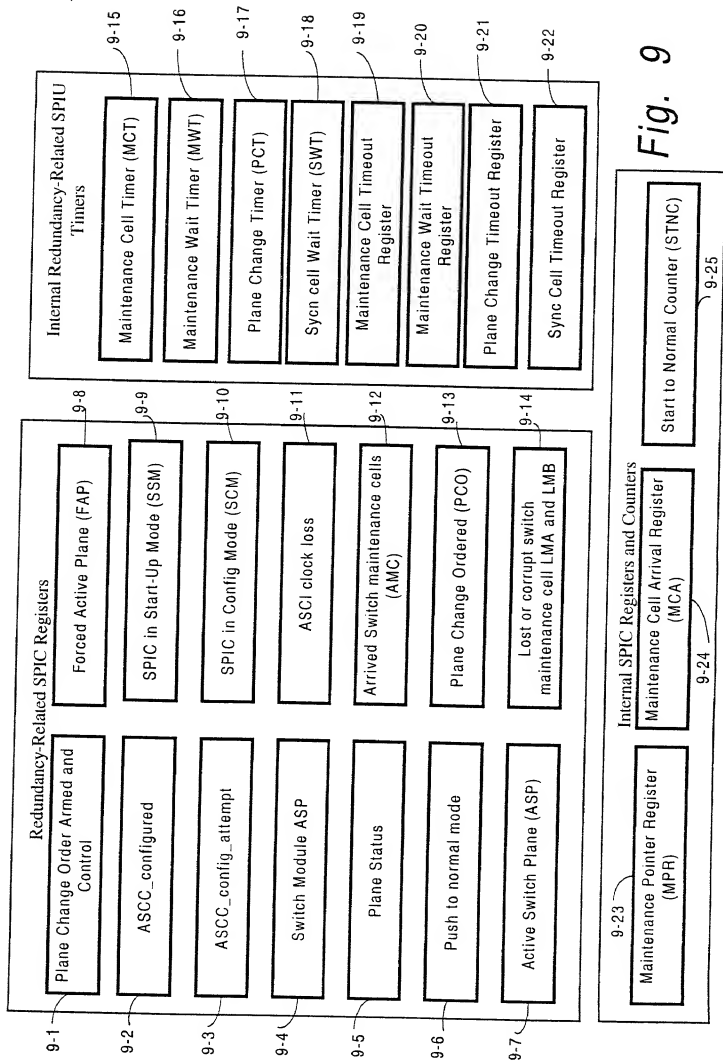


Fig. 9

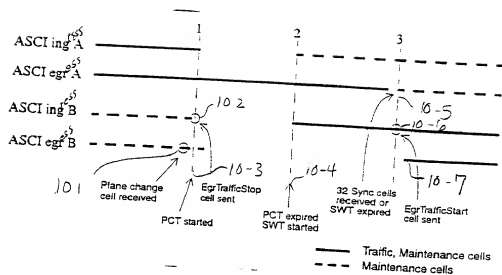
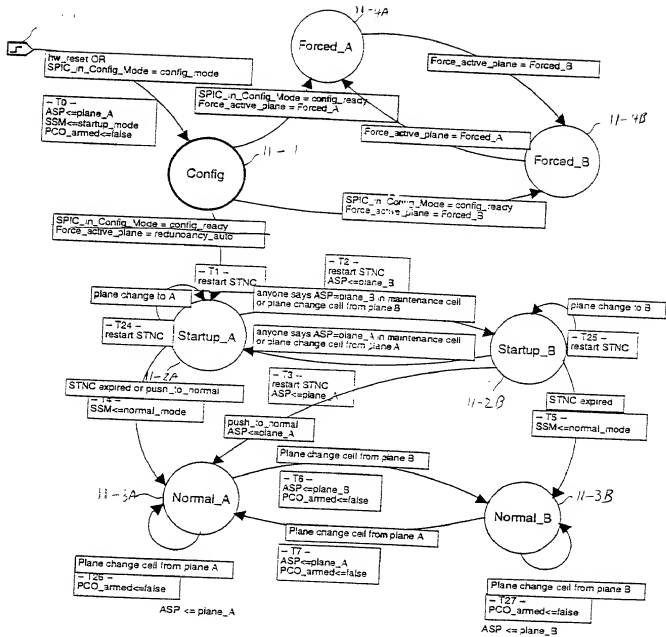


Fig 10



ASP = Active switch plane
 SSM = SPIC n startup mode
 PCO_armed = Plain change order armed
 STNC = Start to normal counter

Fig 11




Fig 12




Fig 13